

CIRCA PREPRINT 2010/9

REWRITING THE CHECK OF 8-REWRITABILITY FOR A_5

ALEXANDER KONOVALOV

ABSTRACT. The group G is called n -rewritable for $n > 1$, if for each sequence of n elements $x_1, x_2, \dots, x_n \in G$ there exists a non-identity permutation $\sigma \in S_n$ such that $x_1 x_2 \cdots x_n = x_{\sigma(1)} x_{\sigma(2)} \cdots x_{\sigma(n)}$. Using computers, Blyth and Robinson (1990) verified that the alternating group A_5 is 8-rewritable. We report on an independent verification of this statement using the computational algebra system GAP, and compare the performance of our sequential and parallel code with the original one.

Let $n > 1$ be an integer. Following [1], a group G is said to be *totally n -rewritable*, or have the rewriting property P_n , if for each sequence of n elements x_1, x_2, \dots, x_n of the group G there exists a non-identity permutation $\sigma \in S_n$ such that

$$x_1 x_2 \cdots x_n = x_{\sigma(1)} x_{\sigma(2)} \cdots x_{\sigma(n)}.$$

Clearly, all abelian groups satisfy P_2 , and if G satisfies P_k then it also satisfies P_{k+1} .

On the problem session of the conference “Arithmetic of Group Rings and Related Objects” (Aachen, Germany, March 22-26, 2010) Eli Aljadeff (Technion, Haifa, Israel) suggested the following problem:

Prove that the alternating group A_5 has the property P_8 .

He referred to the computer verification of this statement reported in [1], and demonstrated how to show that A_5 has the property P_{10} using group rings technique. He also suggested that since [1] appeared twenty years ago, nowadays this result probably could be verified much faster. Motivated by this, the author verified that A_5 has the property P_8 using the computational algebra system GAP [2] and compared the performance of the sequential and parallel GAP implementations with the one described in [1].

To check that the group G is n -rewritable using the brute force approach, one may enumerate all n -tuples of distinct elements of G and check that each of them may be rewritten. Of course, even for A_5 the number of tuples to check will be enormous, so this approach will not work.

There is, however, a simple observation that allows to reduce the number of checks substantially. The algorithm described by Blyth and Robinson in [1] constructs all non-rewritable words of length 2 which are non-equivalent with respect to the action of $\text{Aut}(G)$. On the next step, these words are used to construct all non-equivalent non-rewritable words of length 3, and then the process is repeated

2000 *Mathematics Subject Classification.* Primary 20B35; Secondary 20B40.
Key words and phrases. Alternating group, rewritability, GAP.

until the first n for which there are no non-rewritable words of length n will be found.

The next table contains information about the number $N(r)$ of non-rewritable words of length r in A_5 , determined using computers and listed in [1]:

r	$N(r)$
2	29
3	1315
4	43121
5	528069
6	187719
7	1320
8	0

The authors of [1] wrote that first these data were computed over a period of two weeks by a PASCAL program on a MicroVAX II, and then verified by a parallel C++ implementation that produced the same result on four Sun 3/60 machines in less than three hours.

We were interested to compare the reported performance with the runtime that can be achieved nowadays in a sequential version of GAP on modern computers. To test our implementation, we used an 8-core Intel server, with dual quad-core Intel Xeon 5570 2.93GHz / RAM 48 GB / CentOS Linux 5.3. First we present the algorithm in the pseudocode, following its textual description from [1]:

```

REWRITABILITYLENGTH( $G, m$ )
1   $A \leftarrow \text{AUTOMORPHISMGROUP}[G]$ 
2   $x \leftarrow \text{NONTRIVIALORBITREPRESENTATIVES}[A, G]$ 
3   $n \leftarrow 1$ 
4  repeat
5       $n \leftarrow n + 1$ 
6       $nrw \leftarrow \text{EMPTY LIST}$ 
7      for  $u$  in  $x$        $\triangleright u$  is a non-rewritable word of length  $n - 1$ 
8          do  $K \leftarrow \text{INTERSECTION}[$ 
                   $\text{STABILISER}[A, u[1]],$ 
                   $\dots,$ 
                   $\text{STABILISER}[A, u[n - 1]]]$ 
9          if  $\text{SIZE}[K] > 1$ 
10         then  $y \leftarrow \text{NONTRIVIALORBITREPRESENTATIVES}[K, G]$ 
11         else  $y \leftarrow G \setminus \{1_G\}$ 
12         for  $v$  in  $y$ 
13             do  $t = \text{CONCATENATION}[u, v]$ 
14                 if  $\text{ISREWITABLEWORD}[t]$ 
15                 then  $\text{APPEND}[nrw, t]$ 
16         if  $nrw = \text{EMPTY LIST}$ 
17             then return  $n$ 
18         else  $x \leftarrow nrw$ 
19     until  $n = m$ 
20 return fail

```

The pseudocode above refers to the following procedures:

- `AUTOMORPHISMGROUP`[G] returns the automorphism group of G ;
- `NONTRIVIALORBITREPRESENTATIVES`[A, G] return the list of representatives of orbits of non-identity elements of the group G under the action of its automorphism group A ;
- `STABILISER`[A, g] returns the stabiliser of an element $g \in G$ in the group A ;
- `IsREWRITABLEWORD`[t] checks if the word t is rewritable.

Other names of procedures should be self-explanatory.

The first implementation looked very much like the pseudocode above, and it took almost 34 hours to run (though it used one CPU, other CPUs were used for other jobs, so we can not guarantee exact measurement). The second version was optimised to achieve more efficiency on the stage when most of non-rewritable words of length k can not be extended to non-rewritable words of length $k + 1$. Concatenation of lists was replaced by changing the last element “in place”, and `IsREWRITABLEWORD` was inserted directly into the loop without a call to a separate function that, in its turn, used `ForAny`. Additionally, intersection of stabilisers was computed in a loop which breaks if a trivial subgroup is constructed. This and some other minor optimisations allowed to reduce the runtime to about 15 hours. Finally, we traded space vs time and stored not only tuples, but also stabilisers of their elements in $\text{Aut}(G)$. This permitted further speedup and reduced the runtime to be less than ten hours. At this stage we performed six clean measurements on a machine not running other user’s jobs, and the average runtime was 9 hours and 41 minute.

The GAP code for the function `RewritabilityLength` is given in the Appendix. As you can see from the example of a GAP session below, the numbers of non-rewritable words exactly coincide with the data from [1]:

```
gap> G := AlternatingGroup(5);;
gap> Exec("date");RewritabilityLength(G,10);time;Exec("date");
Wed Mar 31 11:04:39 BST 2010
Started enumeration of NRW of length 2
29 NRW of length 2 constructed
Started enumeration of NRW of length 3
1315 NRW of length 3 constructed
Started enumeration of NRW of length 4
43121 NRW of length 4 constructed
Started enumeration of NRW of length 5
528069 NRW of length 5 constructed
Started enumeration of NRW of length 6
187719 NRW of length 6 constructed
Started enumeration of NRW of length 7
1320 NRW of length 7 constructed
Started enumeration of NRW of length 8
0 NRW of length 8 constructed
8
33383583
Wed Mar 31 20:21:09 BST 2010
```

Furthermore, it is easy to see that we can process independently each non-rewritable word of the length k to derive all non-rewritable words of length bigger

than k . Clearly, this allows parallelisation. The parallel version of the algorithm was implemented using the master-worker skeleton from the GAP package SCSCP [3]. To ensure that no data were lost, the output was modified to return the total number of non-rewritable words of each length summing the numbers over all parallel procedure calls:

```
gap> RewritabilityParallel(AlternatingGroup(5),10,4);
[ 0, 29, 1315, 43121, 528069, 187719, 1320, 0 ]
```

The third parameter specifies that parallel computation will be started from the words of length four. Thus, in the beginning 1315 non-rewritable words of length are computed sequentially to ensure an optimal task granularity.

The program was tested first on an 8-core Intel server, with dual quad-core Intel Xeon 5570 2.93GHz / RAM 48 GB / CentOS Linux 5.3 with 1 master and 8 workers, and then on a cluster consisting of three machines of the same configuration as above with 1 master and 24 workers. The average runtime on six measurements is given in the table below.

Number of workers	Runtime	Speedup	Efficiency = speedup/nr.workers
4	3h 0m 5s	3.23	0.81
8	1h 50m 23s	5.26	0.66
16	55m 42s	10.43	0.65
24	37m 16s	15.59	0.65

To summarise, we have provided an independent verification of the result from [1]. This may be considered as an additional motivation to find a theoretical proof that A_5 is 8-rewritable. While the original publication 20 years mentions the usage of PASCAL or C++ for sequential and parallel computations respectively, now this computation has been implemented in GAP, being compatible with other GAP code and readable by a suitably qualified GAP user. Parallel tools offered in the SCSCP package made it possible to parallelise the code within the GAP system without the necessity to switch to other traditional for the high-performance computing languages that support parallelism. Note that an ongoing HPC-GAP project (<http://www-circa.mcs.st-and.ac.uk/hpcgap.php>) is aimed to reengineer the GAP system to provide better support for shared and distributed memory programming models, so in the future this example may be hopefully even better reimplemented in a new version of the GAP system.

REFERENCES

- [1] R. Blyth, D. Robinson, Solution of the solubility problem for rewritable groups. *J. London Math. Soc.* (2) 41 (1990), no. 3, 438–444.
- [2] The GAP Group, *GAP – Groups, Algorithms, and Programming, Version 4.4.12*; 2008, <http://www.gap-system.org>.
- [3] A. Konovalov, S. Linton. SCSCP — Symbolic Computation Software Composability Protocol, Version 1.2; 2010, <http://www.cs.st-andrews.ac.uk/~alexk/scscp.htm>.

APPENDIX: GAP SOURCE CODE FOR THE SEQUENTIAL VERSION

```

RewritabilityLength:=function(G,limit)
local eltsG, s, A, orbsA, x, q, n, isnrw,
      nrw, S, i, j, u, K, y, orbsK, v, tw;
eltsG:=Filtered( G, s -> s <> () );
A:=AutomorphismGroup(G);
orbsA:=Orbits(A,G);
x:=Filtered( List(orbsA, q -> q[1] ), q -> q <> () );
x:=List( x, q -> [ [ q ], Stabilizer( A,q ) ] );
n:=1;
repeat
  n:=n+1; nrw:=[];
  Print("Started enumeration of NRW of length ", n, "\n");
  S := SymmetricGroup( n );
  S := Filtered( S, s -> s <> () );
  for i in [1..Length(x)] do
    # Print( i, "/", Length(nrw), "\r");
    u := x[i][1]; K := x[i][2];
    if Size(K) = 1 then
      y := eltsG;
    else
      orbsK:=Orbits(K,G);
      y:=Filtered( List(orbsK, q -> q[1] ), q -> q <> () );
    fi;
    tw := u;
    for v in y do
      tw[n] := v;
      isnrw:=true;
      for s in S do
        if Product(tw)=Product(Permuted(tw,s)) then
          isnrw := false; break;
        fi;
      od;
      if isnrw then
        Add( nrw, [ ShallowCopy(tw),
                    Intersection( K, Stabilizer( A,v ) ) ] );
      fi;
    od;
    Print( Length(nrw), " NRW of length ", n, " constructed\n");
    if nrw=[] then return n; fi;
    x := ShallowCopy( nrw );
  until n=limit;
  return fail;
end;

```

SCHOOL OF COMPUTER SCIENCE, UNIVERSITY OF ST ANDREWS, ST ANDREWS, FIFE, KY16 9SX
SCOTLAND

E-mail address: alexk@mcs.st-andrews.ac.uk